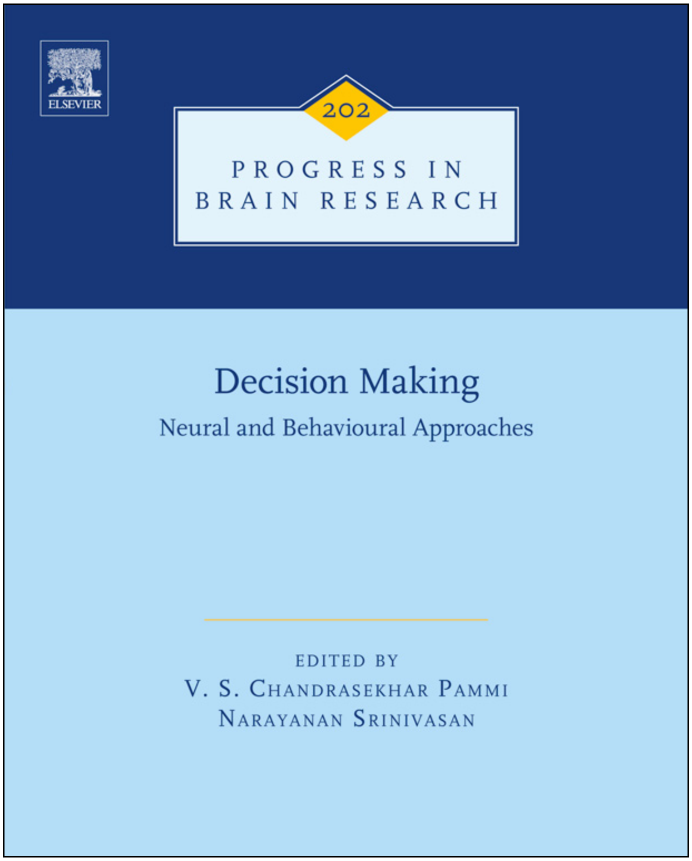

From: Adrian R. Camilleri and Ben R. Newell, Mind the gap? Description, experience, and the continuum of uncertainty in risky choice.
In V.S. Chandrasekhar Pammi and Narayanan Srinivasan, editors:
Progress in Brain Research, Vol. 202,
Amsterdam: The Netherlands, 2013, pp. 55-71.
ISBN: 978-0-444-62604-2

CHAPTER

Mind the gap? Description, experience, and the continuum of uncertainty in risky choice

4

Adrian R. Camilleri[*], Ben R. Newell[†,1]

[*]*Fuqua School of Business, Duke University, Durham, NC, USA*

[†]*School of Psychology, University of New South Wales, Sydney, Australia*

[1]*Corresponding author. Tel. +61-2-93851606, Fax: +61-2-9385 3641, e-mail address: ben.newell@unsw.edu.au*

Abstract

The description–experience "gap" refers to the observation that choices are influenced by whether information about potential alternatives is learnt from a summary description or from the experience of sequentially sampling individual outcomes. In this chapter, we traverse the cognitive steps required to make a decision—information acquisition, storage, representation, and then choice—and at each step briefly review the evidence for sources of discrepancy between these two formats of choice. We conclude that description- and experience-based choice formats lie along a continuum of uncertainty and share important core features, including the explicit representation of probability, the combining of this probability information with outcome information, and utility maximization. The implication of this conclusion is that the differences between description- and experience-based choices emerge from how uncertainty information is acquired and stored rather than how it is represented or used.

Keywords

description-experience gap, sampling bias, recency, uncertainty, probability, exemplar model

On February 1st, 2003 the space shuttle *Columbia* disintegrated over Texas and Louisiana during its reentry into the Earth's atmosphere. Tragically, all seven crew members aboard perished in the disaster. Later, investigations revealed that the cause of the accident was a breach in the thermal protection system owing to damage sustained during launch when a piece of insulation foam broke off and hit the leading edge of the left wing (NASA, 2008). The disaster sparked intense debate about the risks associated with space flight and the very future of NASA space missions.

Progress in Brain Research, Volume 202, ISSN 0079-6123, http://dx.doi.org/10.1016/B978-0-444-62604-2.00004-6

From the perspective of a cognitive psychologist, it is thought provoking to consider both the risk information and the format in which it was conveyed to the NASA personnel prior to their decision to participate in the doomed *Columbia* flight. The crew members had access to two formats of risk information. The first format of risk information was their own previous experience observing past flights. At the time, the mission was the 113th space shuttle launch and *Columbia's* 28th mission. During that time, only one disaster had previously occurred when the space shuttle *Challenger* broke apart in 1986 and killed all seven crew members. The second format of risk information was the risk estimate described by NASA engineers. Based on information gathered from the *Challenger* accident and other near misses, NASA had computed the probability of losing a shuttle and its crew to be about 1% per flight (Buchbinder, 1989).

In this particular case, the two information formats—previous experience and explicit descriptions—provide very similar risk information. As a result, it might then appear straightforward to conclude that the information format the NASA personnel relied upon to make their choice—in this case to participate in the mission—was inconsequential. Interestingly, however, the results of several recent experimental studies cast doubt over this intuition. In this chapter, we review the literature contrasting decisions from experience with decisions from description and then draw some conclusions about where these two formats appear to truly produce different choices. To frame the discussion, we traverse the cognitive steps required to make a decision—information acquisition, storage, representation, and then choice—and at each step briefly review the evidence for sources of discrepancy between these two formats of choice. We conclude that experience- and description-based choice formats lie along a continuum of uncertainty and can indeed produce different choices, but also share important core features, including the explicit representation of probability, the combining of this probability information with outcome information, and utility maximization.

1 WHAT IS THE DESCRIPTION–EXPERIENCE CHOICE "GAP"?

A "decision from experience" is defined as a choice situation in which the alternative decision outcomes and their associated probabilities are learned from observing a sequential sample of outcomes over time. Referring back to the introductory example, evaluating the risk of space flight disaster by observing the outcome of previous space shuttle launches would qualify as an experience-based choice. In contrast, a "decision from description" is defined as a choice situation in which the alternative decision outcomes and their associated probabilities are learned from a summary description explicitly stating this information. Referring back to the introductory example once again, evaluating the risk of space flight disaster by reading the executive summary of NASA's 1989 risk analysis report would qualify as a description-based choice. The distinction between description- and experience-based choices has become of particular interest in the past few years because of substantial evidence demonstrating that preferences tend to systematically diverge depending on which information format is relied upon—this phenomenon has since been termed

Table 1 Comparison of the different choice paradigms

Choice Paradigm	**Format**	**Description**	**Experience**		
			Sampling	**Partial Feedback**[a]	**Full Feedback**[a]
	Graphical depiction[b]	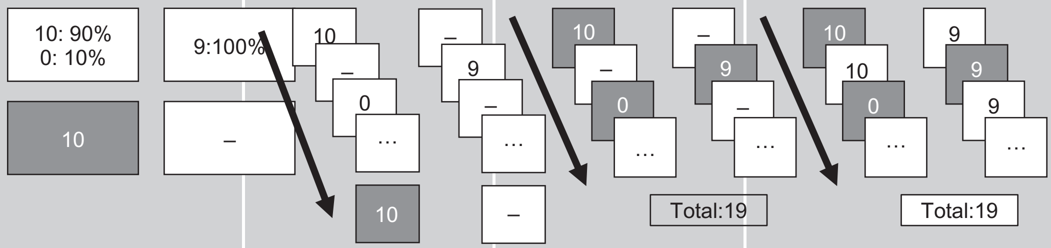			
Key characteristics	Outcome distribution	Known	Unknown	Unknown	Unknown
	Choice type	Single	Single	Repeated	Repeated
	Feedback type	Incomplete	Incomplete	Incomplete	Complete
Illustrative problems (% selecting the R [isky] option)[c]	*S*: 9 *R*:10(0.9)[d]	15[e–g]	38[g]	60	70
	S: −3[d] *R*: −4(0.80)	58[f,g]	40[f]	15	20
	S: 2[d] *R*: 14(0.15)	53[f]	38[f]	5[g]	30
	S: −3 *R*: −32(0.10)[d]	45[g]	48[g]	65	80

[a] *The DV was the choice made on the final (i.e., 100th) trial.*
[b] *Shaded rectangles represent consequential trials, that is, trials in which the outcome of the choice affected earnings.*
[c] *S, safe option; R, risky option. Data originally reported in Camilleri and Newell (2011c).*
[d] *Option predicted to be preferred if rare events are underweighted.*
[e] *Significantly different from Sampling condition ($\chi^2 < 0.05$).*
[f] *Significantly different from Partial Feedback condition ($\chi^2 < 0.05$).*
[g] *Significantly different from Full Feedback condition ($\chi^2 < 0.05$).*

the "description–experience gap" and can be thought of as assignment of more psychological weight to rare events when described than when experienced (Hertwig and Erev, 2009; Rakow and Newell, 2010).

Experience-based choices have primarily been studied using the three different paradigms graphically represented in the top-most part of Table 1. In the *Partial Feedback* paradigm, the decision-maker is presented with the alternative options and encouraged to sample outcomes from each option in any order. Each sample briefly reveals a randomly selected outcome, with replacement, from a hidden distribution associated with the option. Crucially, each sampled outcome adds to a running total that is constantly displayed to the decision-maker. The decision-maker is not informed how many samples will be granted but is encouraged to earn the highest score. Thus, the decision-maker is faced with a tension between the objectives of learning more about the options ("explore") while also trying to maximize earnings across an unknown number of repeated, consequential choices ("exploit"; Cohen et al., 2007). Surprisingly, Barron and Erev (2003) observed that participants in the Partial Feedback group showed opposite patterns of choice to participants in the Description group: certain outcomes were less attractive rather than more attractive, risk aversion was displayed in the loss domain rather than in the gain domain, and decisions were made as if rare events were underweighted rather than overweighted. The exploration–exploitation tension inherent to the Partial Feedback paradigm can be mitigated by also providing feedback for the foregone alternative. This *Full Feedback* paradigm has been shown to produce experience-based preferences that also appear to underweight rare events (e.g., Yechiam and Busemeyer, 2006).

The exploration–exploitation can also be eliminated by separating these competing goals into distinct phases, which is the rationale behind the *Sampling* paradigm. During the initial sampling phase, the decision-maker is encouraged to sample outcomes from each option in any order. Importantly, each sampled outcome during this phase is without financial consequence and is purely for the purpose of learning the outcome distribution associated with the option. At any point during the sampling phase, the decision-maker can elect to stop sampling and move on to the choice phase. During the choice phase, the decision-maker selects the option that he/she prefers with the goal of earning the highest score. Using this paradigm, Hertwig et al. (2004) observed large choice differences depending on whether participants were learning about the outcome distributions in description or experience formats.

The three experience paradigms outlined above share many features in common, mostly notably permitting the decision-maker to sequentially experience a series of outcomes. Moreover, the pattern of preferences between the different experience conditions is similar: For example, there is a very strong, positive correlation between preferences observed with the Partial Feedback paradigm (Barron and Erev, 2003) and the Sampling paradigm (Hertwig et al., 2004). There also appears to be a close correspondence between the paradigms in the alternation rate between the available options that diminishes as the number of trials used increases (Gonzalez and Dutt, 2011).

Many studies have now found evidence consistent with the idea that rare events seem to be given more weight when described than when experienced, which has the effect of producing a description–experience choice gap (see Hertwig and Erev, 2009; Rakow and Newell, 2010). Although we have pointed out the similarities between the three experience tasks, there are also some critical differences in terms of the number of choices and type of feedback that we thought might also be important upon close inspection (see the middle section of Table 1). We decided to carefully examine these differences in a recent investigation (Camilleri and Newell, 2011c). To facilitate comparisons, the experience-based paradigms were equated in terms of the number of trials, problems, and instructions. The contrast between the Sampling and Partial Feedback conditions was important to discover the influence of making repeated choices. The contrast between the Partial and Full Feedback conditions was important to discover the influence of the exploration–exploitation tension. As shown in the bottom-most of Table 1, we replicated the basic description–experience choice gap. More importantly, we found a large difference between Sampling and two Feedback conditions, but no difference within the Feedback conditions (i.e., between the Partial and Full Feedback conditions).[1] These observations are crucial to understanding the mechanisms contributing to the gap, which is a discussion we now turn to.

2 WHAT ARE THE CAUSES OF THE DESCRIPTION–EXPERIENCE CHOICE "GAP"?

There are several potential causes of the description–experience gap of which some have been investigated in more depth than others. We frame the discussion within the conceptual framework presented in Fig. 1, which attempts to isolate each potential stage between acquiring information and making a choice. Note that the framework summarized in Fig. 1 represents a convenient scaffold from which to launch our discussion rather than a strict endorsement.

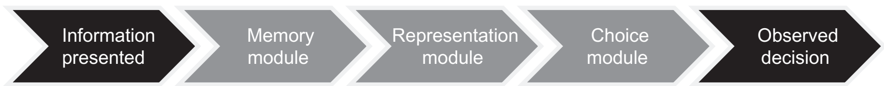

FIGURE 1

A conceptual framework incorporating the potential stages at which description- and experience-based decisions might diverge. Black chevrons represent external, observable events. Gray chevrons represent internal, mental events.

[1]There was a tendency for participants to select the riskier option more often in Full Feedback condition, which is consistent with previous studies that show a hot stove effect: less risky choices when feedback is limited to the chosen option (Erev et al., 2010). In this chapter, we do not examine this phenomenon further but see Erev and Haruvy (in press) for more information.

2.1 Differences in acquired information?

The first stage in making a choice in an uncertain environment is to gather information. In a description-based decision, information acquisition is easy and accurate. By contrast, in an experience-based decision information acquisition can be difficult and biased because sequentially sampling outcomes from a static distribution does not ensure that the observed sample will be representative of the underlying distribution (Hertwig et al., 2004). This issue of misleading, or biased, samples is particularly important in the sampling paradigm where small samples are often taken. Such small samples, when taken from a skewed binomial distribution, can be shown to result in fewer encounters with the rare event than expected from the objective probability (Hertwig and Pleskac, 2010). For example, if 1000 people each draw 20 samples from an option containing a rare outcome with an objective probability of 0.1, just 28.5% will encounter the rare event as expected. In contrast, 32.3% of people will see the rare outcome more than expected and the majority of people—39.2%—will experience the rare event less than expected, if at all. This threat of misleading samples is particularly relevant in the sampling paradigm because participants often display very frugal sampling behaviors and usually take a median of just 5–10 samples per option (Hau et al., 2010). Such frugal sampling is thought to make choices easier by amplifying the differences between options (Hertwig and Pleskac, 2008). Consistent with this hypothesis, Hertwig et al. (2004) found that 78% of participants had sampled the rare event less than expected, and this experience had a distinct impact on choices. For example, in the fourth example shown in Table 1—a sure loss of 3 versus a 10% chance of losing 32—only 46% of participants preferred the risky option when the rare loss of 32 was encountered as frequently as or more frequently than expected. In contrast, *all* participants preferred the risky option when the rare loss of 32 was encountered less frequently than expected.

Subsequent research has debated whether the description–experience gap can be entirely explained as a statistical phenomenon caused by misleading samples. Fox and Hadar (2006) conducted a reanalysis of the Hertwig et al. (2004) data and found that Prospect Theory (Kahneman and Tversky, 1979) could satisfactorily account for both description- and experience-based choices when based on the outcome probabilities actually experienced by the participants (as opposed to the objective, underlying outcome probabilities). Also in support of the statistical account, Rakow et al. (2008) yoked the description-based problems faced by one group of participants to the actual outcome distributions observed by another group of participants facing experience-based problems. They found that elimination of misleading samples also eliminated the choice gap. However, Hau et al. (2010) subsequently showed that this null effect was carried predominately by cases in which samples had been particularly frugal and had rendered the choice trivial (e.g., 100% chance of $3 vs. 100% chance of $0). In a strictly controlled study examining this issue, Camilleri and Newell (2011a) eliminated the possibility of misleading samples by allowing participants the freedom to select the number of perfectly representative sample sets to observe. We found that under these conditions the choice gap was all but eliminated.

Other studies have observed the choice gap even in the absence of misleading samples. Ungemach et al. (2009) removed the impact of sampling bias by obliging participants to sample 40 times from each option while ensuring that all samples were representative of the underlying outcome distribution. For example, a participant faced with problem described above would eventually select the risky options 40 times and observe \$32 exactly 4 times and \$0 exactly 36 times. Participants were free to sample the options in any order, and the order of the outcomes was random. They found that although the size of the gap was reduced when compared to those in a free sampling condition, it was not eliminated. This finding was supported by three other studies in which participants observed a large number of samples either by providing large incentives (Hau et al., 2008, Experiment 1) or simply by obliging a large sample (Camilleri and Newell, 2011c; Hau et al., 2008). As shown in the columns of Table 1 comparing the Description and Sampling conditions, although the choice gap closed in size, it nevertheless remained apparent when averaging across problems in the Camilleri and Newell (2011c) data.

Together, these results suggest that decision-makers' choices are often the same regardless of whether examined in the description or sampling paradigm when equivalent information is relied upon. However, the story clearly does not end here. As is obvious from Table 1, there are cases where the gap is observed even in the presence of large samples that closely match the underlying distribution (i.e., the feedback paradigm). Thus, additional explanatory mechanisms further along the conceptual framework shown in Fig. 1 are clearly required.

2.2 Differences in how acquired information is stored?

Once information has been acquired, it must be stored in memory in some manner (Fig. 1). Differences between description and experience formats may arise if different types of information are stored. Moreover, the sequential nature of the experience-based choice format additionally allows for the potential influence of memory order effects.

In general, there are two broad storage system types that have been considered: exemplar and nonexemplar. An exemplar-type system explicitly represents and stores each outcome that is observed. The Instance-based Learning (IBL) model (Lejarraga et al., 2012) is an example of a successful choice model with an exemplar-type memory system: the model compares and then selects the alternative with the highest "blended value," which is the summation of all observed outcomes weighted by their probability of retrieval. Importantly, each observed outcome is individually stored as an "instance" along with other contextual information. In contrast, a nonexemplar-type system does *not* explicitly represent or store each particular unit of information but instead combines each observed outcome in some way and then only stores the combined element. The value-updating model (Hertwig et al., 2006) is an example of a choice model with a nonexemplar-type memory system: the model calculates the value of an option as the weighted average of the previously estimated value and the value of the most recently experienced outcome. Importantly, each observed outcome is discarded and only the updated value is stored.

The format of description-based choices has ensured that models designed to account for such decisions nearly universally incorporate an exemplar-type memory system that explicitly records outcome information (see Brandstatter et al., 2006, for a review). In contrast, models designed to account for experience-based choices have shown greater variability in storage type. A review of the literature, however, reveals that exemplar-type models have performed better in all recent experience-based model competitions (Erev et al., 2010; Gonzalez and Dutt, 2011; Hau et al., 2008) and also hold additional explanatory potential (e.g., to account for inaccurate probability estimates, see below).

As described earlier, sequentially observing a sample that is representative of the underlying distribution does not ensure that all outcomes will be weighted equally, or even considered, when making a choice. Such potential memory order effects are particularly relevant given that research on memory (Atkinson and Shiffrin, 1968) and belief updating (Hogarth and Einhorn, 1992) have demonstrated that the order in which outcomes are experienced can influence the weight accorded to those outcomes. Moreover, according to Kareev's (1995, 2000) narrow window hypothesis, people tend to make inferences based on a limited number of items in working memory, and hence, decisions are often based on a subset of experiences. Memory order effects could contribute to the choice gap if later sampled outcomes are weighted more heavily than earlier sampled outcomes because rare events are less likely than common events to have occurred recently and thus less likely to affect choice.

In support of the importance of memory order effects, Hertwig et al. (2004) found that the second half of sampled outcomes did indeed predict choices better than the first half of sampled outcomes (75% vs. 59%, respectively). Thus, participants demonstrated a recency effect whereby outcomes observed later in the sequence were given relatively more weight when making the choice. We observed a recency effect in the data shown in Table 1, which we used to explain the small choice gap remaining between the Description and Sampling conditions. However, other experiments have produced mixed support for recency as a contributor to the choice gap: Rakow et al. (2008) found a recency effect for participants in an active sampling condition but not for those in a passive sampling condition, Rakow and Rahim (2010) found a recency effect for children but the opposite effect for adults. In addition, the description–experience gap has been observed in absence of memory order effects (Camilleri and Newell, 2011a; Hau et al., 2008; Ungemach et al., 2009) and in cases without memory burden at all (Hau et al., 2010).

Together, these results suggest that memory order effects, especially in the form of recency, can contribute to the choice gap but is not a primary cause. Although both exemplar- and nonexemplar-type systems can account for memory order effects by adding weighting parameters, we see greater promise in models that incorporate exemplar-type memory storage systems.[2]

[2]For exemplar-type models, the frequency and similarity of stored outcomes is also crucial although these issues have not been addressed as extensively in the experience-based choice literature (Gonzalez and Dutt, 2011; Nosofsky, 1988).

2.3 Differences in how probability information is represented in the mind?

Information storage and representation are clearly intimately connected; nonetheless, we believe that the two can be discussed separately because distinct causes of the gap could occur either during storage or in representation. For example, another potential source of difference between description and experience formats is how probability information is represented in the decision maker's mind: one format may explicitly represent probability information whereas the other may not. In the case where both formats explicitly represent probability information in the mind, the gap could still emerge if decision-makers systematically misrepresent probability information as function of information format. Indeed, although frequency information appears to be automatically stored (Hasher and Zacks, 1984), estimates of probability can often be inaccurate (Erev et al., 1994; Lichtenstein et al., 1978; Zacks and Hasher, 2002) and even the same information presented in physically different formats can be represented and subsequently used quite differently (Gigerenzer and Hoffrage, 1995).

Although the debate continues, the inference from the description-based choice literature appears to be that probability information is indeed explicitly represented. This conclusion stems from the finding that choice models that explicitly represent probability information better predict choices than models that do not. For example, the minimax strategy, which simply selects the option with largest experienced minimum outcome, and other choice heuristics that ignore probability information have been shown to have limited success in predicting description-based choices (Brandstatter et al., 2006). In contrast, the most successful models in the description-based choice field have been those that explicitly represent probability information, in particular, "weighted utility" models such as cumulative prospect theory (Tversky and Kahneman, 1992; Tversky et al., 2004) and its variants (Erev et al., 2010).

The debate is even livelier in the experience-based choice literature. In particular, the natural mean heuristic (Hertwig and Pleskac, 2008), which simply selects the option that produces the largest average outcome during sampling, has shown an impressive ability to predict choices (Hau et al., 2008). Similarly, traditional reinforcement learning models assume that only the running average is stored and no representation of probability information is retained. In spite of the general appeal of such simple models, experience-based choice models that ignore probability information are generally out-performed by those models that do not (Erev et al., 2010; Gonzalez and Dutt, 2011). A further classification can be made with respect to models that do not ignore probability information: those that implicitly store outcome probabilities and those that explicitly store outcome probabilities. An example of implicitly stored probability information is the IBL model described earlier: each past experience is recorded in terms of context, choice, and outcome; given this information outcome probabilities can be computed but are not explicitly stored (Gonzalez and Dutt, 2011). In contrast, an example of explicitly stored probability information is the two stage model (Fox and Tversky, 1998).

The relative performance of the different choice models suggests that probability information is unlikely to be entirely ignored. Several behavioral experiments have followed-upon this assertion by directly asking decision-makers to provide estimates of outcome probability. As noted by Fox and Hadar (2006), the gap could be explained if probabilities are differentially estimated as a function of information format (e.g., overestimated in the description format but underestimated in the experience format). We directly tested this possibility by asking participants to provide estimates of the probability of each outcome in several gamble problems (Camilleri and Newell, 2009). To allow for the possibility that decision-makers do not numerically represent probabilities when options are learned from experience, judgment probes were either verbal (i.e., asked to enter a number) or nonverbal (i.e., asked to adjust the density of a grid). Consistent with past research (e.g., Barron and Yechiam, 2009; Gottlieb et al., 2007; Ungemach et al., 2009), we found that rare events were consistently overestimated and, promisingly, more so in the description condition (which was also replicated in Camilleri and Newell, 2011b). However, there was no evidence that the effect of presentation format on choice was mediated by its effect on probability estimates.

Together, these results suggest that probability information may be explicitly represented in the mind in both description and experience formats, and, based on this representation, decision-makers tend to overestimate the probability of rare events. However, there is little evidence that this misrepresentation of probability information is a cause of the gap. Before moving on from this section it is worthwhile highlighting a phenomenon that might be called the "overestimating–underweighting paradox": the observation made in the context of experience-based choice that people tend to *over*estimate rare events yet behave as if they *under*weight them (Barron and Yechiam, 2009; Marchiori et al., submitted).

In one study where the overestimating–underweighting paradox was observed, we took participants through a 3-stage test procedure: First, learn about the alternative outcome distributions either from description or via sampling; second, enter a probability estimate corresponding to how often the rare event occurs in each alternative distribution (note that those in the description condition simply had to retain the description information in short-term memory); and third, indicate the preferred option (Camilleri and Newell, 2011b). Problem 1, for example, was a choice between 3 for sure and an 80% chance of 4 (and 20% chance of 0, which is the rare event). The risky option was preferred by 36% of participants in the Description group (consistent with overweighting of the rare event) but 64% of participants in the Experience group (consistent with underweighting of the rare event). However, probability estimates measured before making the choice found that participants estimated that the rare event occurred 35% of the time in the Description group (an overestimation) and 27% of the time in the Experience group (also an overestimation). It is clear that any complete model of experience-based choice must account for this puzzle; one potential candidate is discussed in the next section.

2.4 Differences in how the representations are contrasted to make a choice?

The final step prior to making a physical choice is to contrast the representations associated with each alternative and apply some sort of decision rule (Fig. 1). Differences between description- and experience-based choices could emerge if different decision rules are employed as a function of information format. Choice rules vary in terms of whether options are valued independently or only in comparison with one another (Vlaev et al., 2011). At one extreme is a "value-first" rule type whereby the decision-maker forms a preference for the option that is independently computed to be associated with the highest value (e.g., Prospect Theory, Kahneman and Tversky, 1979). At the other extreme is a "comparison-only" rule type whereby the decision-maker forms a preference through direct comparison of the available options, potentially without the calculation of value at all (e.g., Priority Heuristic, Brandstatter et al., 2006).

The most successful decision rule in the description-based choice literature is a weighted utility rule (Chen and Corter, 2006). The utility of an option is calculated as the sum of each value multiplied by its probability of occurring, with some weightings applied. In Prospect Theory (Kahneman and Tversky, 1979), for example, the value and probability weighting functions are nonlinear: the value function implies diminishing sensitivity to increases in the absolute payoffs from the reference point, and the probability weight function implies that decision-makers overweight low probabilities and underweight moderate and high probabilities. According to the rule, the alternative that promises the highest utility is preferred. Thus, the weighted utility rule is of a "value-first" rule type.

In the experience-based choice literature, there are a number of different choice rules that vary in complexity and success. As noted earlier, the choice rule of the natural mean heuristic is to simply select the alternative that has produced the highest mean outcome. A much more complex choice rule is employed by the "ensemble" model, which assumes that each choice is made based on the average prediction of four equally likely rules: two versions of the k-sampler model, a version of stochastic cumulative prospect theory, and a version stochastic priority heuristic (Erev et al., 2010). In the middle range of complexity is the choice rule of the IBL model, which selects the alternative with the highest "blended value", which is calculated as the summation of all observed outcomes weighted by their probability of retrieval.

Isolating the "best" choice rule is difficult given that it is a nonobservable process that follows from previous nonobservable processes (see Fig. 1). In addition, vastly different choice rules have enjoyed some success. For that reason, we suggest three criteria for endorsing a choice rule: coherence, parsimony, and predictive validity. First, in order to cohere with memory and representational processes, the rule should be based on an exemplar-type memory system that explicitly stores outcome and probability information. Second, the rule should be as simple as possible and minimize the need to introduce additional free parameters. If possible, the rule should be broadly applicable. The most obvious candidate is a simple utility rule similar to the one endorsed by most description-based choice models. Third, the choice rule should successfully predict choice behavior, at least relative to alternative rules.

We recently proposed an experience-based choice model designed to meet these criteria: the exemplar confusion (ExCON) model (Hawkins et al., in preparation). The goal of the ExCON model is to provide an account of how sequentially sampled outcomes are used to form a representation of the outcome distribution and how that representation is used to form a preference. The model is broadly aimed and designed to account for both probability estimates and choices in the Feedback and Sampling paradigms. The ExCON is an exemplar-based model that stores each observed outcome on every trial. Similar to the IBL model described earlier, the ExCON implicitly represents probability information via its storage of individual exemplars. Crucially, the storage of each exemplar is associated with a small probability of memory interference such that currently stored exemplars can become "confused." The memory store is envisioned to be limitless and all stored exemplars—veridical or otherwise—are equally considered at the point of choice. The ExCON choice rule is of a "value-first" type that combines each outcome with its estimated probability of occurring and then selects the option that maximizes utility.

In order to rigorously test the ExCON model, we conducted an experiment that presented participants with binary choices between five-outcome options in the Sampling and Feedback paradigms (Hawkins et al., in preparation). We also asked each participant to estimate the probability associated with each outcome. The ExCON model was able to account for the tendency to overestimate rare outcomes and also did well at predicting choice preferences, which also showed a tendency to underweight rare events.

When the ExCON model was entered into the Technion Prediction Tournament (Erev et al., 2010), it won the Sampling competition and came close to winning the Feedback competition. Thus, the model appears to be a very strong candidate. Perhaps more importantly, its simple utility decision rule is directly imported from existing models of description-based choice, suggesting that the decision rule may not be a source of difference between the description and experience formats of choice.

3 WHERE DO WE GO FROM HERE?

We'll never know what format of risk information the NASA personnel relied upon prior to participating in the doomed Columbia flight. What is clear, however, is that reliance on personal experience often causes us to form a preference that is different to the one we would have formed if presented with the true outcome distributions. A key reason for this description–experience choice gap can be attributed to a reliance on inaccurate representations of the world. In most cases, our experiences are very limited and so decisions are made based on a relatively small sample of outcomes. A small sample of outcomes frequently misrepresents the true distribution of outcomes in the world, most often under-representing rare events. This external sampling bias is often combined with an internal sampling bias. The internal sampling bias can be most readily attributed to a noisy memory system that may rely on more recently sampled outcomes. Such reliance often compounds the under-representation of rare events in the sample relied upon to make a choice and produces preferences that are consistent with underweighting of rare events. Our review reveals that in decision contexts where a single choice is made subsequent to learning about the options (i.e., the Sampling paradigm), then the difference between description and experience choice formats can be reduced when a representative sample is used as the basis of choice (i.e., when external and internal sampling biases are eliminated; e.g., Camilleri and Newell, 2011a).

Of course, experience-based choices rarely occur in a vacuum after a lengthy period of costless sampling and reflection. Instead, we usually make experience-based choices on the fly and while simultaneously learning more about the outcome distributions associated with the alternative options. In such situations when each sampled outcome is consequential, preferences can still be consistent with underweighting of rare events even when samples are perfectly representative of the world (i.e., the Feedback paradigm). The difference between these experiential tasks—that is, costless sampling followed by a choice (the sampling paradigm) and repeated consequential sampling (the feedback paradigm)—does not appear to be attributable to the tension between the goals of exploring and exploiting the options in the latter format because the difference remains even in the context of complete feedback. Our review therefore reveals that the difference between description and experience choice formats is also attributable to the sequential nature of the experience-based choices.

Samples of outcomes acquired sequentially must be combined in some way to represent the outcome distribution. Decision-makers tend to overestimate rare outcomes and underestimate more common outcomes when asked to explicitly report outcome distributions or to nonverbally represent them. Thus, people do not appear to perfectly weigh and combine sequentially observed outcomes. We suggest that decision maker's judgment inaccuracies reflect the processes of a noisy memory system. This system is embodied in ExCON model. The model also shows that obtained probability estimates are only useful in predicting choices when combined with a utility function implying diminishing marginal utility. Our review therefore reveals that explicit probability representation is an important feature of experience-based choice and that another key difference between description- and experience-based choices is how probabilistic information is stored—in experience-based choice, this process appears to be based on noisy, instance-based memory.

Rather than conceptualizing description- and experience-based choices as discrete, we prefer to represent them as lying along a continuum of uncertainty (Hau et al., 2010; Rakow and Newell, 2010; cf. Knight, 1921). There are two observations that support this continuum of uncertainty argument. First, when the unique features of experience-based choice are eliminated, then preferences often become the same as those observed in the description format. The unique features of experience-based choices are the need to search the environment for information and the need to repeatedly integrate this information into a representation. These unique features give rise to the sources of difference between description and experience: sequential sampling of outcomes, acquisition of biased samples of information, and reliance on noisy memory. Crucially, when these differences are accounted for—by eliminating the sequential nature of the choice, by presenting representative samples, and by manipulating the sequence of outcomes to be cyclical—then choice differences disappear.

Second, the models that best account for experience- and description-based choices explicitly represent probability information and share a common choice mechanism. Based on the results from the Technion Prediction Tournament, description-based choices are best modeled with a stochastic version of cumulative prospect theory (SCPT). Our most recent work suggests that experience-based choices are best modeled with the ExCon model (Hawkins et al., in preparation). The SCPT and ExCon models both explicitly represent probability information, combine this with outcome information, and then maximize utility as suggested by axioms of rationality (Bernoulli, 1738/1967).

The notion of a continuum is in contrast to proposals suggesting that description- and experience-based choices are conceptually unique and therefore require fundamentally different theories of choice. Accordingly, models of choice that do not at least implicitly represent probability and combine it with outcome information—including choice heuristics and reinforcement models—fail to completely capture the psychological mechanisms involved in experience-based choice.

If decisions under uncertainty do lie along a common continuum, then the primary goal of future research is to produce a single, complete model of choice under uncertainty. Such a model would simultaneously account for experience- and description-based choices. The scaffolding used in this review, and the success of the ExCon model in particular, demonstrates the potential value of separately conceptualizing and then bolting together different basic cognitive processes to produce complex processes like those that occur when making decisions under uncertainty. With this analogy as inspiration, a complete model of choice under uncertainty would be constructed from basic components that are combined and activated under different choice conditions. From the perspective of experience-based choice, more work is required to improve understanding of the search component (e.g., Hills and Hertwig, 2010). From the perspective of description-based choice, more work is required to improve understanding of how descriptions of probability are represented in the mind (e.g., Gottlieb et al., 2007).

The insights provided here are not limited to the theoretical. Beyond the walls of the lab individuals, organizations, and governments continually rely on experience to guide decisions under uncertainty. Research into experience-based choice may help to explain why rare events such as the 1993 attack on New York's World Trade Center or the 1988 savings and loan crisis often fail to adequately alter behavior or policy to reduce the likelihood of future unwanted "black swan" events (Taleb, 2007). The findings also help to explain why different people may hold conflicting opinions about important social issues such as nuclear energy use, immunization, or the need to act on climate change despite having access to ostensibly equivalent information (Weber, 2006). Ultimately, the best choices will be made by those of us who recognize the limitations inherent in our information and memory capability, and seek out information from all points along the description–experience continuum.

References

Atkinson, R.C., Shiffrin, R.M., 1968. Human memory: a proposed system and its control processes. In: Spence, K.W. (Ed.), The Psychology of Learning and Motivation: Advances in Research and Theory. Academic Press, New York, pp. 89–195.

Barron, G., Erev, I., 2003. Small feedback-based decisions and their limited correspondence to description-based decisions. J. Behav. Dec. Making 16, 215–233. doi:http://dx.doi.org/10.1002/bdm.443.

Barron, G., Yechiam, E., 2009. The coexistence of overestimation and underweighting of rare events and the contingent recency effect. Judgm. Decis. Making 4, 447–460.

Bernoulli, D., 1738. Exposition of a New Theory on the Measurement of Risk. Gregg Press, Farnborough Hants, England (L. Sommer, Trans.).

Brandstatter, E., Gigerenzer, G., Hertwig, R., 2006. The priority heuristic: making choices without trade-offs. Psychol. Rev. 113, 409–432.

Buchbinder, B., 1989. Independent Assessment of Shuttle Accident Scenario Probabilities for the Galileo Mission, vol. 1. NASA/HQ Code QS, Washington DC.

Camilleri, A.R., Newell, B.R., 2009. The role of representation in experience-based choice. Judgm. Decis. Making 4, 518–529.

Camilleri, A.R., Newell, B.R., 2011a. Description- and experience-based choice: does equivalent information equal equivalent choice? Acta Psychol. 136, 276–284.

Camilleri, A.R., Newell, B.R., 2011b. The relevance of a probabilistic mindset in risky choice. In: Carlson, L., Hölscher, C., Shipley, T. (Eds.), Proceedings of the 33rd Annual Conference of the Cognitive Science Society. Cognitive Science Society, Austin, TX, pp. 2794–2799.

Camilleri, A.R., Newell, B.R., 2011c. When and why rare events are underweighted: a direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. Psychon. Bull. Rev. 18, 377–384.

Chen, Y.J., Corter, J.E., 2006. When mixed options are preferred in multiple-trial decisions. J. Behav. Dec. Making 19, 17–42.

Cohen, J.D., McClure, S.M., Yu, A.J., 2007. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. Phil. Trans. R. Soc. 362, 933–942.

Erev, I., Haruvy, E., in press. Learning and the economics of small decisions. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics, vol. 2.

Erev, I., Wallsten, T.S., Budescu, D.V., 1994. Simultaneous over- and underconfidence: the role of error in judgment processes. Psychol. Rev. 101, 519–527.

Erev, I., Ert, E., Roth, A.E., Haruvy, E.E., Herzog, S., Hau, R., Lebiere, C., 2010. A choice prediction competition, for choices from experience and from description. J. Behav. Dec. Making 23, 15–47.

Fox, C.R., Hadar, L., 2006. "Decisions from experience" = sampling error + prospect theory: reconsidering Hertwig, Barron, Weber and Erev (2004). Judgm. Decis. Making 1, 159–161.

Fox, C.R., Tversky, A., 1998. A belief-based account of decision under uncertainty. Management Sci. 44, 879–895.

Gigerenzer, G., Hoffrage, U., 1995. How to improve Bayesian reasoning without instruction: frequency formats. Psychol. Rev. 102, 684–704.

Gonzalez, C., Dutt, V., 2011. Instance-based learning: integrating sampling and repeated decisions from experience. Psychol. Rev. 118, 523–551.

Gottlieb, D.A., Weiss, T., Chapman, G.B., 2007. The format in which uncertainty information is presented affects decision biases. Psychol. Sci. 18, 240–246.

Hasher, L., Zacks, R.T., 1984. Automatic processing of fundamental information: the case of frequency of occurrence. Am. Psychol. 39, 1372–1388.

Hau, R., Pleskac, T.J., Kiefer, J., Hertwig, R., 2008. The description-experience gap in risky choice: the role of sample size and experienced probabilities. J. Behav. Dec. Making 21, 1–26.

Hau, R., Pleskac, T.J., Hertwig, R., 2010. Decisions from experience and statistical probabilities: why they trigger different choices than a priori probabilities. J. Behav. Dec. Making 23, 48–68.

Hawkins, G., Camilleri, A.R., Newell, B.R., Brown, S.D., in preparation. The substituted exemplar model: a model of estimation and choice in a sequential sampling paradigm.

Hertwig, R., Erev, I., 2009. The description–experience gap in risky choice. Trends Cogn. Sci. 13, 517–523.

Hertwig, R., Pleskac, T.J., 2008. The game of life: how small samples render choice simpler. In: Charter, N., Oaksford, M. (Eds.), The Probabilistic Mind: Prospects for Bayesian Cognitive Science. Oxford University Press, Oxford, England, pp. 209–236.

Hertwig, R., Pleskac, T.J., 2010. Decisions from experience: why small samples? Cognition 115, 225–237.

Hertwig, R., Barron, G., Weber, E.U., Erev, I., 2004. Decisions from experience and the effect of rare events in risky choice. Psychol. Sci. 15, 534–539.

Hertwig, R., Barron, G., Weber, E.U., Erev, I., 2006. The role of information sampling in risky choice. In: Fiedler, K., Juslin, P. (Eds.), Information Sampling and Adaptive Cognition. Cambridge University Press, New York, pp. 75–91.

Hills, T.T., Hertwig, R., 2010. Information search in decisions from experience: do our patterns of sampling foreshadow our decisions? Psychol. Sci. 21, 1787–1792.

Hogarth, R.M., Einhorn, H.J., 1992. Order effects in belief updating: the belief-adjustment model. Cogn. Psychol. 24, 1–55.

Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica 47, 263–291.

Kareev, Y., 1995. Through a narrow window: working memory capacity and the detection of covariation. Cognition 56, 263–269.

Kareev, Y., 2000. Seven (indeed, plus or minus two) and the detection of correlations. Psychol. Rev. 107, 397–402.

Knight, F.H., 1921. Risk, Uncertainty, and Profit. Sentry Press, New York.

Lejarraga, T., Dutt, V., Gonzalez, C., 2012. Instance-based learning: A general model of repeated binary choice. J. Behav. Dec. Making 25, 143–153.

Lichtenstein, S., Slovic, P., Fischhoff, B., Combs, B., 1978. Judged frequency of lethal events. J. Exp. Psychol. 4, 551–578.

Marchiori, D., Guida, S.D., Erev, I., submitted. The four-fold reaction to rare events and the role of small samples.
NASA, 2008. Columbia crew survival investigation report.
Nosofsky, R.M., 1988. Similarity, frequency, and category representations. J. Exp. Psychol. Learn. Mem. Cogn. 14, 54–65.
Rakow, T., Newell, B.R., 2010. Degrees of uncertainty: an overview and framework for future research on experience-based choice. J. Behav. Dec. Making 23, 1–14.
Rakow, T., Rahim, S.B., 2010. Developmental insights into experience-based decision making. J. Behav. Dec. Making 23, 69–82.
Rakow, T., Demes, K.A., Newell, B.R., 2008. Biased samples not mode of presentation: re-examining the apparent underweighting of rare events in experience-based choice. Organ. Behav. Hum. Decis. Process. 106, 168–179.
Taleb, N.N., 2007. The Black Swan: The Impact of the Highly Improbable. Random House, New York.
Tversky, A., Kahneman, D., 1992. Advances in prospect theory: cumulative representation of uncertainty. J. Risk Uncertainty 5, 297–323.
Tversky, A., Kahneman, D., Shafir, E., 2004. Advances in Prospect Theory: Cumulative Representation of Uncertainty Preference, Belief, and Similarity: Selected Writings by Amos Tversky. MIT Press, Cambridge pp. 673–702.
Ungemach, C., Chater, N., Stewart, N., 2009. Are probabilities overweighted or underweighted, when rare outcomes are experienced (rarely)? Psychol. Sci. 20, 473–479.
Vlaev, I., Chater, N., Stewart, N., Brown, G.D.A., 2011. Does the brain calculate value? Trends in Cognitive Sciences 15, 546–554.
Weber, E.U., 2006. Experience-based and description-based perceptions of long-term risk: why global warming does not scare us (yet). Clim. Chang. 77, 103–120.
Yechiam, E., Busemeyer, J.R., 2006. The effect of foregone payoffs on underweighting small probability events. J. Behav. Dec. Making 19, 1–16.
Zacks, R.T., Hasher, L., 2002. Frequency processing: a twenty-five year perspective. In: Sedlmeier, P., Tilmann, B. (Eds.), Frequency Processing and Cognition. Oxford University Press, New York, NY, pp. 21–36.